Designing a Michelangelo-Class Humanoid Sculptor: Specifications for Embodied Technical and Aesthetic Mastery

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

December 11, 2025

What would it actually take to build a humanoid robot capable of sculpting marble with something like Michelangelo's fluency, not only copying forms but inventing new ones in a recognizably related style? This paper treats that question as a concrete design problem rather than a loose metaphor. It proposes a layered specification for a "Michelangelo-class" sculpting robot that integrates high-force, high-precision mechanics; rich visual, tactile, and force sensing; compliant, impact-capable control; and an AI "art stack" trained on the geometry, tool marks, and compositional language of Renaissance sculpture. Rather than imagining a disembodied generative model that outputs meshes for others to fabricate, the focus here is on an embodied agent that works in stone, with chisels, rasps, and mallets, under conditions of dust, vibration, and fracture risk. The specification is organized into a body stack, skill stack, and art stack, and then unfolded into an end-to-end workflow from initial prompt or maquette to finished statue. Along the way, the paper proposes evaluation metrics, Turing-style tests for sculptural mastery, and broader implications for craft, authorship, and the future of mixed human–robot ateliers.

Keywords: humanoid robotics, robotic sculpture, Michelangelo, dexterous manipulation, aesthetic evaluation, embodied intelligence, art and AI, style modeling, human–robot collaboration, creative automation.

A collaboration with GPT-5.1-Thinking-Extended. 49 pages. CC4.0.

Outline

1. Introduction
1.1 Framing the question: from metaphor to specification
1.2 Why Michelangelo is a meaningful benchmark
1.3 Three stacks for a sculpting robot: body, skill, and art
1.4 Scope, assumptions, and intended audience

2. Historical and Aesthetic Context
2.1 Michelangelo's tools, materials, and studio practice
2.2 Non-finito and the aesthetics of emergence from stone
2.3 Anatomy, contrapposto, and the drama of gesture
2.4 What it means to be "akin to Michelangelo" rather than identical

3. Capability Targets and Comparisons to Existing Systems
3.1 Material and tool requirements: marble, maquettes, and more
3.2 Fidelity goals: geometry, surface, and tool-mark signatures
3.3 Originality versus reproduction: style-consistent invention
3.4 Limitations of current CNC carving and industrial robotic arms

4. The Body Stack: Mechanics and Sensing
4.1 Morphology and work envelope of a sculpting humanoid
4.2 Arms, hands, and actuation for high-force, fine-control strikes
4.3 Vision and 3D perception for large blocks and subtle surfaces
4.4 Tactile, force, and vibration sensing in tools and fingertips
4.5 Environmental sensing in a dusty, dynamic atelier

5. The Skill Stack: Control, Manipulation, and Tool Use
5.1 Whole-body motion planning around a massive stone block
5.2 Compliant, impact-capable control loops for chiseling
5.3 Modeling stone, chips, and tool dynamics under uncertainty
5.4 Phase-structured workflows: roughing, blocking, and finishing
5.5 Safety, fail-safes, and human–robot interaction in the studio

6. The Art Stack: Style, Aesthetics, and Creative Judgment
6.1 Data foundations: scans, photographs, and art-historical annotation
6.2 A "Michelangelo-Net" for generative and discriminative style modeling
6.3 Multi-scale aesthetic critics: from local tool marks to global pose
6.4 Balancing style fidelity, diversity, and structural sanity
6.5 Interfaces for human critique, correction, and co-authorship

7. End-to-End Workflow from Prompt or Maquette to Statue
7.1 Input modes: copying, guided design, and creative prompts
7.2 Planning roughing passes with structural safety margins
7.3 Iterative scan–evaluate–carve cycles and convergence criteria
7.4 Handling non-finito and controlled emergence as design choices
7.5 Examples of hypothetical projects and process traces

8. Metrics, Benchmarks, and Turing Tests for Sculptural Mastery
8.1 Perceptual indistinguishability tests with experts and lay viewers
8.2 Geometric, curvature, and surface-statistics comparisons
8.3 Structural robustness and long-term stability of finished works
8.4 Assessing originality within a learned stylistic manifold
8.5 Toward standardized benchmarks for robotic fine-art practice

9. Broader Implications for Craft, Labor, and Ateliers
9.1 Shifting roles of human sculptors in a robot-populated studio
9.2 New pedagogical uses: robotic assistants as didactic tools
9.3 Archival and conservation applications of sculpting robots
9.4 Extensions to other crafts: wood, metal, and architectural ornament
9.5 Ethical questions around authorship, value, and cultural heritage

10. Conclusion and Future Directions
10.1 Summary of the proposed Michelangelo-class specification
10.2 Technological gaps and research priorities
10.3 From single sculptor-robots to mixed-intelligence ateliers
10.4 Open questions at the boundary of engineering and aesthetics

11. References

1. Introduction

Building a humanoid robot that can sculpt “akin to Michelangelo” is an intentionally provocative target. It forces robotics, AI, and aesthetics into the same frame: not just “can a robot carve stone,” but “can a machine inhabit a craft tradition so deeply that its work stands beside canonical human art?” The question is often used metaphorically, as a shorthand for “extreme creative ability,” yet for design and research it is more fruitful to treat it literally. What sensors, actuators, control strategies, and models of style would actually be required for such a machine to walk into a workshop, pick up a chisel, and bring a human figure out of stone in a way that feels continuous with a Renaissance master?

This paper approaches that challenge as a concrete specification problem. It does not claim that such a system is imminent, nor that it would exhaust the mystery of artistic genius if built. Instead, it uses the Michelangelo benchmark as a way to push beyond generic talk of “creative AI” and “dexterous robots,” compelling us to specify what kind of embodiment, practice, and judgment are needed when the medium is marble and the reference class is one of the most demanding in art history.

1.1 Framing the question: from metaphor to specification

Most discussions of artistic AI remain at the level of metaphor: a model “paints like Van Gogh,” a robot arm “sculpts like a master.” In reality, the underlying systems are usually narrow: they either generate images in a particular style, or they follow precomputed tool paths from a digital mesh. The gap between these demonstrations and a studio-trained sculptor is enormous.

To move from metaphor to specification, the question must be rephrased in operational terms: what observable capabilities, under what conditions, would justify saying that a humanoid robot sculpts akin to Michelangelo? This demands attention to at least four aspects:

- The range of materials and tools it must handle.
- The precision and repeatability of its movements under load.
- The nature of its internal representations of form, pose, and style.
- The ways in which it evaluates its own work and decides what to do next.

The rest of the paper takes this reframed question as its organizing principle. Rather than arguing abstractly about whether machines can be “truly creative,” it specifies what a Michelangelo-class sculpting robot would need to look like and do, and how one might evaluate its outputs.

1.2 Why Michelangelo is a meaningful benchmark

Michelangelo is not chosen simply because he is famous or canonical. His practice concentrates, in a particularly demanding form, many of the challenges that make this design problem interesting.

First, he worked in marble at monumental scales. This requires an understanding of structural integrity, fracture behavior, and the limits of slender limbs and dynamic poses. A system that can only produce small, fragile pieces or that frequently fractures the stone is not operating in the same domain.

Second, Michelangelo's figures are anatomically acute and dynamically composed. Contrapposto, torsion, and complex weight shifts are not incidental stylistic flourishes; they encode a deep understanding of the human body in motion and at rest. Capturing this kind of gestural intelligence in a generative model is far harder than matching superficial surface textures.

Third, his surfaces carry distinctive tool marks and transitions. The interplay between rough and polished, plane and curve, finished and deliberately unfinished (non-finito) forms part of his artistic language. A robot that merely delivers smooth, neutral surfaces, or that copies a mesh without regard to tool-trace signatures, misses an important dimension of the benchmark.

Finally, Michelangelo sits at the intersection of technical mastery and highly legible public significance. Works like David or the Pietà are not obscure; they have been scrutinized by generations of viewers and scholars. Any serious claim that a robot is achieving something "akin to Michelangelo" will inevitably be tested against demanding, well-informed critics, which is precisely what makes the benchmark useful.

1.3 Three stacks for a sculpting robot: body, skill, and art

To make the problem tractable, this paper organizes the specification into three interacting stacks:

- The body stack concerns the robot's physical form and sensors. It includes morphology, joints, actuators, grippers, and the vision, depth, force, tactile, and vibration sensing necessary to work safely and precisely on large stone blocks. It asks: what must the machine be physically capable of doing and perceiving?
- The skill stack captures low- and mid-level control policies. It covers whole-body motion planning, compliant and impact-capable control, tool dynamics, and phase-structured workflows for roughing, blocking, and finishing. It asks: how does the machine translate high-level sculpting intentions into sequences of safe, effective actions in a messy, uncertain environment?

- The art stack encodes style, aesthetic judgment, and design. It includes data curation, style modeling, multi-scale critics that evaluate intermediate states, and interfaces for human input. It asks: how does the robot know what counts as a good Michelangelo-like pose, surface, or degree of non-finito, and how does it balance fidelity with originality?

Throughout the paper, these stacks are treated as modular for clarity, but they are meant to work together in closed loops. Embodied artistic practice arises from their interaction, not from any one layer alone.

1.4 Scope, assumptions, and intended audience

This paper is speculative but engineering-oriented. It does not present an existing system; rather, it proposes a target architecture and capability profile that could guide future work in humanoid robotics, dexterous manipulation, and artistic AI. Several simplifying assumptions are made for clarity:

- The robot operates primarily in a dedicated studio space with known infrastructure.
- It can rely on substantial off-board computation and storage if needed.
- Training data, including high-resolution scans and images of canonical works, are assumed to be available.
- Questions of legal rights, conservation ethics, and institutional adoption are deferred to later sections.

The intended audience is interdisciplinary. Roboticists may find in the body and skill stacks a set of concrete challenges and benchmarks beyond industrial assembly. Researchers in computer vision and generative models may see in the art stack a demanding test case for style modeling that must survive physical instantiation in stone. Art historians and practicing sculptors may use the specification as a framework for articulating what, precisely, makes a sculptural practice coherent and recognizable across works.

By making the Michelangelo question as precise as possible, the introduction sets up the rest of the paper to explore how far contemporary and near-future technologies could go toward answering it, and what that would mean for both engineering and art.

2. Historical and Aesthetic Context

Any serious attempt to specify a Michelangelo-class sculpting robot must begin with the historical realities of Michelangelo’s practice. The challenge is not only to remove stone accurately, but to participate in a way of working where tools, materials, anatomy, and spiritual

symbolism are tightly intertwined. Marble blocks came from particular quarries with known flaws and grain. Tools were limited but refined through generations of craft. Studio practice relied on assistants, models, sketches, and bozzetti, yet the final decisions at the block were intensely personal. The robot's specification must therefore be anchored in this ecology rather than in an abstract notion of "style."

This section sketches four key aspects of that context. First, it reviews the tools, materials, and workshop routines that shaped Michelangelo's daily work. Second, it examines the role of non-finito, the deliberate or circumstantial unfinishedness that has become central to interpretations of his sculpture. Third, it considers anatomy, contrapposto, and the staged drama of gesture as structural features of his figures rather than decorative choices. Finally, it clarifies what it might mean for a robot to sculpt "akin to" Michelangelo without being a perfect imitator or a counterfeit.

2.1 Michelangelo's tools, materials, and studio practice

Michelangelo's primary medium for large-scale sculpture was Carrara marble, quarried in massive blocks that could contain veins, inclusions, and hidden fractures. Working these blocks required a sequence of tools used in a relatively stable progression. Roughing-out often began with pitching tools and point chisels, driven by iron or wooden mallets to knock off large chunks and establish the basic silhouette. Subsequent passes relied on tooth chisels such as the gradina and subia to model forms with ridged grooves, followed by flat chisels, rasps, and rifflers to refine transitions and begin smoothing. Drills could be used to open deep undercuts or accent details like hair and pupils, and abrasives completed the polished areas.

The physical logic of this tool sequence matters for a robot specification. Each stage implies different force regimes, stroke lengths, and tolerances for error. Point chisels can remove several millimeters or more of material per strike but at higher risk of unintended cracking. Tooth chisels can plane and model forms while still leaving clear tool marks. Rasps and rifflers act more like sanding, shifting the task from impact control to pressure and sliding. A robot that simply treats "stone removal" as a uniform process will miss the economic and aesthetic logic embedded in this traditional hierarchy of tools.

Michelangelo's studio practice also involved planning and iteration before touching the final marble. He used drawings, wax or clay models, and sometimes full-scale cartoons or enlarged bozzetti. Assistants could help with squaring up blocks, rough transfers from models, and polishing, but accounts stress Michelangelo's personal control over critical phases, including the most delicate passages. Quarried blocks were sometimes reused or re-envisioned, as in the case of the flawed block that became David. For a robot, this suggests that the system needs not only

low-level carving routines but also support for multi-stage, model-based workflows and the ability to revise plans in light of unexpected flaws discovered during carving.

2.2 Non-finito and the aesthetics of emergence from stone

Michelangelo's unfinished works, especially the series commonly known as the Slaves or Prisoners, have become paradigmatic examples of non-finito: figures that seem to strain and twist their way out of uncarved marble. The Awakening Slave, the Young Slave, the Bearded Slave, and the Atlas Slave all juxtapose roughly hewn stone with partially articulated bodies whose limbs and torsos vanish back into the block. Whether left incomplete for pragmatic reasons or as deliberate aesthetic choices, these works have been read as metaphors of the soul struggling against matter, of latent forms imprisoned in stone, and of artistic creation as an ongoing, never fully finished liberation.

Non-finito is not merely a failure to complete work; it constitutes a recognizable sculptural technique and aesthetic category. The block retains large planar surfaces, rough point-chisel scars, and abrupt transitions where the emerging figure meets untouched marble. The contrast between carved and uncarved zones directs attention to the process itself, making the viewer newly aware that the figure could have remained latent forever. Historical discussions link this to broader Renaissance and neo-Platonic ideas that no earthly work can fully embody its ideal form, and that leaving something "not finished" can be an homage to that incommensurability.

For a Michelangelo-class robot, non-finito has two implications. First, the system must be capable of stopping in aesthetically meaningful states rather than always pushing toward maximal completion. It must be able to treat unfinishedness as a parameter of design, choosing where to leave roughness and where to polish. Second, its internal representations should allow for the figure to be conceptualized as a volume nested in the block, with plans that can leave portions permanently embedded. This is a different design problem than simply approximating a fully detached, free-standing mesh.

2.3 Anatomy, contrapposto, and the drama of gesture

Michelangelo's figures are celebrated for their anatomy and their sense of poised, sometimes explosive, motion. David, for example, stands in a classical contrapposto stance: weight shifted onto one leg, the other relaxed; pelvis and rib cage tilted in opposite directions; shoulders and hips forming counter-angles that give the torso an S-curve. The muscles and tendons are carved with acute sensitivity to tension and relaxation, veins stand lightly under the skin, and even the compression of flesh over bone is carefully observed. This level of physiological detail reflects study of live models, dissection, and the broader Renaissance fascination with the human body as a microcosm.

Beyond anatomical correctness, Michelangelo uses pose to stage psychological and narrative drama. The Dying Slave and Rebellious Slave twist and arch in ways that communicate struggle and surrender. The Slaves in Florence seem to wrestle against their stone prisons, shoulders hunched or torsos straining, the diagonals of their limbs composing a sense of vertical and rotational force. In the Pietà, subtle gestures like the angle of Mary's head or the orientation of her open hand modulate the emotional temperature of the scene, balancing grief with acceptance and calm.

Translating this to a robotic specification means treating anatomy and gesture as structural constraints, not ornament. A generative model must encode not just surface appearance but the volumetric logic of bones, muscles, and skin under gravity. Contrapposto and more complex twists should emerge from underlying articulations rather than being pasted on as arbitrary poses. Furthermore, the system's evaluators must be sensitive to whether a pose reads as tense or relaxed, balanced or on the verge of motion, and whether those qualities align with the intended subject and mood. Without this, any mechanical competence in carving risks producing inert, doll-like figures.

2.4 What it means to be "akin to Michelangelo" rather than identical

Finally, it is important to clarify what sort of benchmark "akin to Michelangelo" represents. Literal indistinguishability from an authentic Michelangelo in every respect is neither necessary nor desirable. Historical works carry patina, context, and authorship histories that no contemporary object, robotic or otherwise, can replicate. Even if a machine could perfectly mimic tool marks and proportions, claiming identity would collapse important distinctions between historic originals and contemporary works.

Instead, "akin to Michelangelo" is best understood as occupying a recognizable region in a stylistic and technical space. A robot meeting this benchmark would produce sculptures that, when viewed without labels, are plausibly grouped by knowledgeable observers with late Renaissance marble figures and, more specifically, with the expressive, anatomically charged, sometimes unfinished works associated with Michelangelo and his immediate circle. Viewers might not mistake them for newly discovered originals, but they would see a clear family resemblance in anatomy, pose, surface handling, and the relationship between figure and block.

From a design perspective, this means aiming for coherence rather than forgery. The art stack should learn the statistical regularities of Michelangelo's work, but it should also be capable of generating novel compositions within that learned manifold. The body and skill stacks should produce tool-use patterns whose traces on the stone are consistent with traditional practice, but not mechanically identical to any known piece. The goal is a robot that can inhabit and

extend a historical sculptural language with integrity, not a machine that counterfeits signatures.

This distinction matters ethically as well. Positioning the robot as a participant in an ongoing craft lineage, rather than as an imitator erasing human authorship, allows for clearer discussions of attribution, value, and collaboration. The benchmark of being “akin to Michelangelo” thus serves as both a technical target and a conceptual boundary, shaping how such a machine might be welcomed or resisted in studios, museums, and public spaces.

3. Capability Targets and Comparisons to Existing Systems

Before specifying actuators, sensors, and models, it is necessary to define what counts as success for a Michelangelo-class sculpting robot. Capability targets anchor the design: they describe the range of materials and tools the system must handle, the level of fidelity expected in geometric and surface outcomes, the kind of originality that is acceptable, and the degree to which current tooling falls short. Without such targets, it is too easy to conflate impressive demonstrations with true parity to a master sculptor.

This section outlines four interlocking capability domains. First, it identifies the material and tool repertoire that a robot must inhabit to be taken seriously as a stone sculptor, including but not limited to marble. Second, it defines fidelity goals in terms of geometry, surface, and tool-mark signatures, recognizing that sculptural language is encoded not just in final form but in the traces of its making. Third, it articulates a distinction between reproduction and style-consistent invention, arguing that simply copying existing works is not sufficient. Finally, it examines the limitations of current CNC carving systems and industrial robot arms, which provide valuable baselines but do not yet meet these artistic demands.

3.1 Material and tool requirements: marble, maquettes, and more

A credible Michelangelo-class robot must work primarily in marble, not only in softer stones or clay. Marble poses a demanding combination of hardness, brittleness, and variability. The grain direction, impurities, and hidden fissures can dramatically change how it responds to impact. A system that can only operate reliably on homogeneous synthetic blocks or foams might be useful for prototyping but would not satisfy the historical and aesthetic benchmark set by Renaissance sculpture.

At the same time, the robot should be fluent in adjacent materials that form part of sculptural practice. Clay and wax maquettes are central to planning and experimenting with forms. A well-equipped sculpting robot ought to be able to work with these softer media, either directly modeling them with hands and small tools or at least scanning and interpreting them as input

prototypes. Plaster casts, wood armatures, and modern materials such as resin might also enter its workflow, especially in mixed-media studios or in conservation contexts.

Tool requirements follow from this material palette. The robot must competently wield a spectrum of tools including point chisels, tooth chisels of various sizes, flat chisels, rasps, rifflers, drills, and abrasives. Each tool has characteristic geometries, weight distributions, and handle shapes that affect grip and motion. The system should be able to pick up, orient, and strike with a mallet in multiple grips, switching tools as needed and returning them to a rack or belt. It must also manage dust and chips with appropriate vacuum or airflow systems, integrating these into its task planning so that visibility and safety are preserved.

Taken together, these requirements imply not only a physically capable body stack but also a flexible tooling and perception infrastructure. The robot is not just a specialized carving head; it is an inhabitant of a workshop in which multiple materials and implements circulate.

3.2 Fidelity goals: geometry, surface, and tool-mark signatures

Fidelity is easily misunderstood as mere geometric accuracy. From a sculptural perspective, however, fidelity has at least three layers: overall geometry, local surface topology, and the microscopic or mesoscopic signatures of tool use.

Geometrically, a Michelangelo-class system should be able to realize forms with sub-millimeter accuracy over human-scale figures. This does not mean that every element must be exact to a predetermined mesh; Michelangelo himself often deviated from models as he discovered opportunities in the stone. Rather, it means that when a target geometry is specified, whether from a scan or a generative model, the robot can reliably approach that geometry, with known error bounds, while respecting the structural constraints of the material.

Surface topology includes smoothness, transitions between planes and curves, and the subtle undulations that give flesh, cloth, and hair their character. Here, the goal is not perfectly uniform surfaces but controlled variation. The system must be able to produce highly polished areas where reflections read cleanly and also maintain regions with deliberate tool marks. It should be capable of matching typical curvature distributions observed in historical works, capturing characteristic bulges, flattenings, and hollows that encode anatomical and stylistic knowledge.

Tool-mark signatures add a further layer. Tooth chisel grooves, rasp textures, and mallet-driven impact patterns leave identifiable traces. A Michelangelo-class robot should be able to reproduce such marks in a way that is statistically similar to traditional practice, including variation in spacing, depth, and orientation. This implies both forward models of how a given stroke will mark the stone and evaluative models that can read back from observed surfaces to determine whether the desired signature has been achieved. Without attention to these

signatures, the results risk looking like generic, smoothed, machine-fabricated objects rather than works rooted in a particular craft lineage.

3.3 Originality versus reproduction: style-consistent invention

There is a fundamental difference between a machine that can copy an existing sculpture and one that can invent new sculptures within a recognizable style. The former is already possible at a basic level: given a high-resolution scan and enough time, current systems can approximate a piece of marble to match the scan within small tolerances. This is an impressive technical achievement but does not meet the conceptual or artistic bar implied by "akin to Michelangelo."

Originality, in this context, does not mean unbounded novelty. Rather, it means the capacity to generate and execute new compositions that belong to the same stylistic family as Michelangelo's works. This includes designing new poses, recombining anatomical motifs, experimenting with degrees of non-finito, and adapting traditional themes to new subjects. The robot should be able to respond to prompts such as "a twisting figure in the spirit of the Slaves, but with a different narrative emphasis," and produce a coherent, physically realizable, and stylistically resonant result.

To support this, the art stack must contain a generative model of the relevant stylistic manifold: a high-dimensional space of forms where local neighborhoods correspond to plausible variations, and global structure reflects deeper constraints like anatomy and dynamics. The model should avoid trivial interpolation between known sculptures and instead learn underlying compositional principles. At the same time, it should include mechanisms for evaluating whether a proposed design is structurally sound in marble and whether it stays within acceptable stylistic bounds. This balances the drive to innovate with the need to remain legible as an extension of a specific tradition.

In practice, this distinction between reproduction and style-consistent invention will be reflected in evaluation protocols. Simply passing geometric similarity thresholds for known works is not enough. The system's outputs must also impress human experts as genuinely new but clearly related, which is a more demanding and ambiguous criterion.

3.4 Limitations of current CNC carving and industrial robotic arms

Existing CNC carving machines and industrial robotic arms provide an important reference point for what is already possible. They can work marble and other stones with high precision, follow complex tool paths, and produce impressive replicas of scanned artworks. However, they fall short of Michelangelo-class capability in several ways that inform the specification of a more advanced system.

First, most CNC setups treat the carving process as a sequence of precomputed passes, with limited adaptation to material variation or emergent aesthetic decisions. If the stone reveals an unexpected vein or fracture, the machine cannot reinterpret the design in situ. It either continues and risks catastrophic failure or stops and waits for human intervention. A Michelangelo-class robot must incorporate real-time sensing and decision-making that can adjust both design and execution as new information about the block emerges.

Second, industrial arms are typically optimized for repeatable, smooth motions under relatively constant loads, not for repeated high-impact strikes with hand tools. While some research systems have demonstrated force-controlled chiseling, mainstream industrial manipulators lack the compliance, backdrivability, and high-frequency control required for nuanced hammer-and-chisel work. They also rarely integrate rich tactile sensing in the gripper or detailed vibration sensing in tools, both of which are crucial for reading the stone's response.

Third, current systems do not embody stylistic understanding. They may carve a Michelangelo statue if given a scan, then switch to a Baroque façade or a contemporary abstract piece, with no internal sense of the differences between these languages. From the machine's perspective, each job is just a different set of tool paths. A Michelangelo-class robot, by contrast, must maintain internal models of form and style that influence how it chooses to work surfaces, where it accentuates or downplays details, and how it balances smoothness against visible tool marks.

Finally, the human–machine interaction patterns around CNC carving are unlike those of a shared atelier. Industrial systems are commonly caged off, programmed offline, and monitored at a distance. The envisioned sculpting robot would instead operate more like a human assistant in a studio: moving around the block, taking direction from a sculptor, and engaging in iterative cycles of critique and revision. This calls for different safety models, user interfaces, and social expectations than those designed for factory floors.

By articulating these limitations, this section clarifies the gap between what present systems can do and the capability targets of a Michelangelo-class humanoid. The following sections build on this gap analysis to specify the body, skill, and art stacks required to close it.

4. The Body Stack: Mechanics and Sensing

The body stack defines what the robot can physically do and perceive before any questions of style or judgment arise. A Michelangelo-class sculptor must be, at minimum, a highly capable industrial machine: it moves heavy tools around a large, irregular workpiece, manages high-impact interactions with brittle stone, and does so for many hours without losing calibration. At the same time, it must approximate human reach, dexterity, and posture options closely enough

that its planned movements map naturally onto sculptural tasks inherited from human practice. Mechanics and sensing therefore have to be designed together: the morphology must make it possible to see and touch the stone from many angles, while the sensor suite must provide enough feedback to make that morphology controllable in a fine-grained way.

This section lays out a conceptual specification for such a body. It begins with the overall form factor and work envelope. It then turns to arms, hands, and actuation, emphasizing the need to combine high peak forces with delicate control. Next, it considers vision and 3D perception for both gross block geometry and subtle surface reading. Finally, it details tactile, force, vibration, and environmental sensing, which together let the robot behave less like a blind machine and more like a practiced craftsperson who can feel and adapt to the stone.

4.1 Morphology and work envelope of a sculpting humanoid

At the scale of Michelangelo's work, the robot's morphology must be matched to large blocks and human-sized figures. A practical design situates the sculptor in the 1.6 to 1.8 meter height range, roughly comparable to an adult human, so that conventional studio ergonomics and tool dimensions remain compatible. Rather than insisting on full bipedal locomotion, a more robust approach uses a wheeled or tracked base with stabilizing outriggers or a retractable pedestal, allowing the robot to lock itself into a rigid stance while carving and then reposition between phases.

The work envelope must accommodate blocks up to several meters tall and wide, with the ability to reach both high and low regions without exotic contortions. This suggests a torso and neck with sufficient pitch and yaw to orient the head and shoulders around the stone, and arms long enough to reach over the top and into moderate undercuts. The body plan should allow the robot to change its effective working height, for example via a vertical lift in the base or adjustable platforms, so that it can adopt comfortable leverage angles for different parts of the figure.

Stability is central. High-energy mallet strikes generate reaction forces and vibrations that propagate through the structure. The base must provide a wide support polygon and a low center of mass, especially when the arms are extended. Outriggers or temporary anchors to the floor can further increase stability during heavy work. Clearances around the block must be considered as well: the robot needs space to swing tools without self-collision, yet remain close enough to maintain precise control. The end result is less a generalized humanoid and more a studio-specific morphology optimized for circling, leaning toward, and working directly against large stone forms.

4.2 Arms, hands, and actuation for high-force, fine-control strikes

The arms are the primary means of delivering energy and precision to the stone. Each arm should have enough degrees of freedom to position and orient tools with human-like flexibility: a three-degree-of-freedom shoulder, a two-degree-of-freedom elbow, and a three-degree-of-freedom wrist provide a useful baseline. This yields redundant kinematics, allowing the robot to adjust joint configurations to maintain favorable leverage and avoid joint limits while carving different regions.

Actuation must support a wide dynamic range. For roughing, the system needs to impart high peak forces through the tool tip, corresponding to firm mallet blows on a point or tooth chisel. For finishing, it must be capable of extremely light, controlled strokes with rasps or small chisels, where a few tenths of a millimeter of unintended motion can visibly damage a delicate surface. Series elastic or variable-stiffness actuators are well suited to this requirement: they allow precise torque control and compliance during impacts, reducing the risk of damaging the stone or the mechanism itself.

Hands and grippers must be designed for sculptural tools rather than generic objects. An anthropomorphic hand with multiple fingers and joints can accommodate a wide variety of grips: hammer grips for mallets, pencil grips for small chisels, lateral pinches for thin rasps, and so on. Fingers should have enough flexion and abduction range to adapt to different handle diameters and shapes. Quick-change couplings at the palm or wrist can further support specialized tool mounts where direct finger gripping is impractical, such as for powered rotary tools.

Crucially, the actuation strategy must consider not just static grip forces but the dynamic coordination between arm and hand during strikes. When delivering a blow with a mallet, the system must time the swing, grip stiffness, and impact angle to produce the desired impulse at the chisel tip while keeping rebound under control. The arms should be capable of modulating stroke length and speed over a continuous range, and of transitioning rapidly between impact strokes and slow, guided movements as the task demands.

4.3 Vision and 3D perception for large blocks and subtle surfaces

Vision is the primary sense through which the robot will perceive form, both at the gross geometric level and at the level of fine surface detail. A layered approach to visual sensing is needed, combining wide-field situational awareness with high-resolution inspection.

At the coarse level, stereo cameras mounted on a head-like structure can provide depth cues and color imagery over a wide field, enabling the system to understand the outline of the block, the position of the figure within it, and the locations of major features. This can be augmented by depth sensors or structured-light projectors to build dense 3D reconstructions of the block

and the current state of the sculpture. The perception system should maintain a consistent coordinate frame tied to the stone, so that measurements made from different vantage points can be fused into a single evolving model.

For subtle surfaces, the robot requires higher resolution than typical depth cameras provide. One option is a dedicated inspection camera with a narrow field of view but high pixel density, mounted on a pan-tilt unit or even on the wrist, allowing close-up views of specific regions. Techniques such as grazing-light inspection, where a controllable light source illuminates the surface at low angles, can exaggerate small bumps and grooves, making tool marks and minor irregularities visible. Photometric methods that infer surface normals from multiple lighting angles can further refine the estimation of curvature and micro-topology.

Because sculpting generates dust and chips, the visual system must be resilient to partial occlusions and transient artifacts. Algorithms should be able to distinguish between stable stone surfaces and temporary debris, and remove the latter from long-term models. Regular re-scanning after sweeping or vacuuming phases can help maintain an accurate representation. The overall requirement is that the robot can “see” the sculpture well enough to compare it to target forms, detect unintended deviations, and plan corrective actions.

4.4 Tactile, force, and vibration sensing in tools and fingertips

While vision provides global context, tactile and force sensing inform the robot about local contact conditions and the stone’s immediate response to actions. Without these, carving reduces to blind execution of planned motions, which is both unsafe and artistically crude.

Force-torque sensors at the wrists can measure the net forces and moments transmitted through tools into the arms. During impact strokes, these sensors can detect whether the chisel has bitten into the stone as expected, whether it has slipped, or whether an unusual spike hints at a crack or inclusion. In more continuous contact tasks, such as rasping or sanding, they allow the system to regulate contact pressure to achieve consistent material removal and avoid gouging.

Tactile sensing in the fingertips and palms adds another layer of information. High-resolution tactile arrays can detect the onset of slip, the distribution of grip forces across a handle, and the presence of small chips or fragments between the hand and the tool. This enables more nuanced and secure tool handling, as well as the possibility of directly touching the stone to feel transitions, edges, and subtle irregularities that may be difficult to see. The robot might, for example, gently run a fingertip along a contour to confirm that a plane is as flat or a curve as continuous as intended.

Vibration sensing, either through dedicated accelerometers on tools or by analyzing signals from existing sensors, is particularly valuable in stonework. The sound and feel of a chisel strike

can reveal whether the stone is solid, whether the tool is sharp, or whether dangerous resonances are building up. A Michelangelo-class robot should be able to classify these signals and adjust its actions accordingly: easing off when chatter indicates a risk of fracture, changing angle when impacts produce dull, unproductive thuds, or stopping entirely if a sudden change suggests that a major crack has opened.

In combination, force, tactile, and vibration sensing allow the body stack to behave less like a rigid machine and more like a responsive instrumented limb, capable of nuanced, feedback-driven interaction with an unpredictable material.

4.5 Environmental sensing in a dusty, dynamic atelier

The studio environment itself presents challenges that the body stack must address. Stone carving generates dust, noise, and flying chips; tools and materials clutter the floor; human collaborators and visitors may move unpredictably around the workspace. Environmental sensing ensures that the robot can operate safely and effectively in this context.

At the most basic level, proximity sensing and obstacle detection are required to prevent collisions. This can be achieved through a combination of depth cameras, lidar, and short-range sensors around the base and torso. The system should maintain a dynamic map of nearby objects, including the block, tool racks, work tables, and people, updating it as conditions change. Safety zones and virtual fences can be defined within this map, constraining the robot's motion planning so that dangerous configurations are automatically avoided.

Dust and air quality monitoring are also important, both for human comfort and for the longevity of the robot's own components. Sensors that measure particle concentrations can trigger ventilation, localized extraction, or temporary pauses in carving. Protective housings and seals on joints and electronics must be designed with the expected dust load in mind, and environmental data can guide maintenance schedules.

Lighting and acoustic sensing further enrich environmental awareness. Knowing the ambient light level and direction helps the vision system calibrate its expectations and control supplemental lights for inspection. Microphones can pick up unusual sounds, such as tool drops or human speech, and route them to higher-level systems for interpretation. Temperature and humidity sensors provide additional context relevant to both stone behavior and equipment performance.

All of these environmental sensing modalities feed into the higher stacks, but they depend on the physical placement and robustness of sensors in the body. A Michelangelo-class sculpting robot is thus not only a pair of skilled arms but a situated agent, continuously reading and responding to the studio around it.

5. The Skill Stack: Control, Manipulation, and Tool Use

If the body stack defines what the robot can potentially do, the skill stack defines what it actually does moment by moment. It is the layer where abstract intentions like “open this plane,” “sharpen that edge,” or “lighten the pressure on this tendon” are translated into trajectories, contact forces, and sequences of tool strokes. For a Michelangelo-class sculptor, this layer must accommodate a wide range of behaviors: circling a massive block without collision, delivering repeated high-energy impacts, sliding tools across the surface with controlled pressure, reacting to unexpected events, and doing all of this in the presence of humans.

Unlike traditional industrial control, the sculpting skill stack must be intrinsically adaptive. The stone is heterogeneous and partially unknown; chips accumulate and then are cleared; the target form may be revised mid-process; human collaborators may intervene with critiques or direct instructions. The control architecture must therefore support closed-loop behaviors at multiple time scales: high-frequency loops for impact and contact, mid-level planners for strokes and local surface shaping, and slower cycles for evaluating progress and revising plans.

This section describes the main components of such a skill stack. It begins with whole-body motion planning around a large stone block. It then examines compliant, impact-capable control loops for chiseling. Next it considers how to model stone, chips, and tools under uncertainty. After that it outlines phase-structured workflows that organize sculpting into roughing, blocking, and finishing. Finally, it discusses safety, fail-safes, and modes of human–robot interaction appropriate to a shared studio.

5.1 Whole-body motion planning around a massive stone block

Working on a monumental sculpture is fundamentally a whole-body task. The robot must position itself relative to the block, orient its torso and head to see the relevant region, and configure its arms to reach the desired point with an appropriate angle and leverage. It has to avoid self-collision, avoid colliding with the stone in unintended ways, and maintain stability while delivering potentially large impulses.

Whole-body motion planning for sculpting differs from planning for assembly or pick-and-place tasks in at least three ways. First, the workpiece is large and irregular, changing shape over time as material is removed. The planner must therefore treat the sculpture as a dynamic obstacle that is also the target of manipulation, updating its internal model as new scans arrive. Second, tool usage requires planning not just end-effector positions but approach and retreat paths that respect the stroke direction and the need for follow-through. A mallet swing, for example, must have enough clear space behind and around the arm to accelerate and decelerate safely. Third,

visibility constraints interact with reach constraints: the robot must sometimes reposition its entire body to get an unobstructed view of the region it intends to work on.

A practical architecture combines global and local planning. Global planners compute safe configurations and paths around the block, including repositioning the base to access new regions. Local planners refine these into concrete arm and hand trajectories that incorporate tool alignment and strike geometry. Throughout, the system must be able to adjust plans in real time when small deviations occur, without losing the higher-level objective. The result is a continuous dance around the stone, where whole-body motions frame and support the fine manipulations of the tools.

5.2 Compliant, impact-capable control loops for chiseling

Chiseling is a distinctive control problem: it involves brief, high-force interactions mediated by a rigid tool, with outcomes that depend sensitively on impact angle, speed, and local material properties. Conventional position or velocity control is not adequate; the robot needs control loops that explicitly manage forces and accommodate the unpredictable aspects of fracture and chip formation.

At the lowest level, the system must regulate the motion of the arm and mallet up to the point of impact, ensuring that the chisel tip is at the intended location and orientation and that the mallet arrives with the planned velocity. Just after impact, it must handle the rapid changes in transmitted forces without saturating sensors or inducing harmful vibrations. Compliance in the joints, whether physical or virtual, allows the mechanism to absorb some of the shock while maintaining enough stiffness to be effective. Control loops at kilohertz rates or higher are appropriate for this regime.

Above that, mid-level controllers translate sculptural intents into sequences of strikes. For a given region, the system may plan a pattern of blows that progressively deepens a groove or lowers a plane. After each strike or small batch of strikes, it uses visual and force feedback to estimate how much material was removed and whether the local stone behavior matched expectations. If the stone is harder or softer than anticipated, or if a strike produces an unexpected fracture, the controller adjusts subsequent strokes: changing their spacing, direction, intensity, or even switching tools.

The same architecture can be adapted to non-impact tasks such as rasping, where contact is continuous rather than impulsive. In that case, the focus shifts to controlling contact force and tangential speed while monitoring friction and vibration signatures. The underlying principle remains: control is always in dialogue with feedback from the material, never assuming that a precomputed trajectory will play out exactly as planned.

5.3 Modeling stone, chips, and tool dynamics under uncertainty

Effective skill requires internal models of how actions change the world. For a sculpting robot, this means models of stone behavior, chip generation, and tool dynamics. These models need not be perfect; in fact, given the complexity and variability of natural stone, exact prediction is impossible. What is needed instead are probabilistic models that capture typical relationships between impact parameters and material removal, and that can be updated on the fly as the robot gains experience with a particular block.

Stone models relate local properties—grain direction, hardness, visible veining—to likely outcomes of different strikes. They inform choices such as whether to work “with” or “against” a vein, how close to cut to a planned surface in a single pass, and how much safety margin to maintain around thin structures. As the robot carves, it can update these models with empirical data: how much depth was achieved per strike, how often micro-cracks appeared, how the sound and vibration signatures correlated with visible results.

Chip and debris models help in planning sequences of strokes that do not bury important details or block sensors. The robot should anticipate where chips will fall, how they will accumulate around the base of the figure, and when they will interfere with footing or visibility. This allows it to interleave carving with clearing actions—vacuuming, sweeping, or blowing chips away—without excessive manual intervention.

Tool dynamics models track the condition of tools over time. A chisel’s edge dulls, changing its cutting efficiency and the forces required to achieve a given effect. Rasps clog with dust, reducing their effectiveness. The robot can use changes in force, vibration, and outcome to infer tool wear and schedule sharpening or replacement. These models also guide tool selection: for a hard region of stone, a sharper tool and a different angle might be preferred; for a more fragile area, a softer approach with a particular rasp could be safer.

All of these models must accommodate uncertainty. Predictions should produce distributions over expected outcomes rather than single values, and control policies should be robust to variation. In effect, the skill stack treats sculpting as a stochastic process, using iterative observation and adaptation to steer it toward desired forms.

5.4 Phase-structured workflows: roughing, blocking, and finishing

Human sculptors typically organize their work into phases: roughing, where large volumes are removed; blocking, where primary forms and planes are established; and finishing, where details and surface qualities are refined. A Michelangelo-class robot should adopt a similar structure, not just for conceptual clarity but because each phase places different demands on control, sensing, and planning.

In the roughing phase, the goal is to approach the general silhouette while staying a safe distance from the intended final surfaces. The robot uses heavy tools, larger strikes, and coarser spatial resolution. Motion planning emphasizes stability and clearance for strong blows. Sensing can tolerate dustier conditions and lower visual precision, as small deviations are unlikely to matter. The internal model tracks how much "stock" remains around the region of interest, ensuring that no overcuts occur that would constrain later decisions.

The blocking phase shifts focus to major planes, volumes, and anatomical landmarks. Tools become somewhat finer, and strikes smaller and more frequent. The robot begins to work with more detailed geometric references: alignments of shoulders and hips, the tilt of the head, the contour of major muscles. Visual inspection and fine-grained scanning become more important, and control tolerances tighten. Non-finito decisions often emerge here: the system may decide, based on aesthetic patterning or instructions, which regions will remain closer to the raw block and which will receive full articulation.

Finally, in the finishing phase, tools are small chisels, rifflers, and abrasives; maneuvers are delicate; and the acceptable error bands narrow further. The robot pays attention to subtle curvature changes, transitions between polished and rough areas, and the continuity of lines and edges. It uses grazing light and close-up inspection to identify minute imperfections, adjusting pressure and stroke direction accordingly. Impact-based actions are minimized or absent, replaced by continuous-contact operations that require precise force regulation.

Phase-structured workflows can be encoded as explicit state machines or as learned policies conditioned on a continuous measure of progress. In either case, the transitions between phases are themselves moments of skill: knowing when to stop roughing, when blocking is "good enough" to support finishing, and when finishing has reached the point where further work would risk over-refinement or loss of desired texture.

5.5 Safety, fail-safes, and human–robot interaction in the studio

The skill stack must be designed with safety and collaboration in mind. Stone sculpting is hazardous even for humans; adding a powerful robot with swinging tools increases the stakes. Safety considerations operate at multiple levels: physical limits on forces and speeds, monitoring for unexpected conditions, and clear interaction patterns with human co-workers.

At the low level, controllers should enforce hard limits on joint torques, end-effector speeds, and the energies of strokes, especially when humans are nearby. Zones of operation can be defined around the robot, with different allowable behaviors depending on whether they are clear or occupied. For example, full-strength roughing might be permitted only when no human is detected within a specified radius, while finishing operations with low forces could be allowed with humans closer.

Fail-safes must handle both internal and external anomalies. Internal anomalies include sensor failures, actuator faults, or software errors; the system should detect these and transition to safe states, such as stopping tool motion, retracting from the stone, and locking the base. External anomalies include sudden human intrusion into an active zone, unexpected movement of the stone, or falling objects. Rapid detection via vision, lidar, or proximity sensors, combined with reflexive braking behaviors, can mitigate risk.

Human–robot interaction protocols determine how sculptors and assistants communicate with the machine. This can involve explicit commands—spoken, gestural, or via interfaces—that start, pause, or redirect work. More subtle interactions might include pointing to regions of the stone, sketching over a 3D model, or annotating scans with desired changes. The skill stack must be able to incorporate such inputs without becoming brittle: it should treat human instructions as constraints or new goals that alter its planning rather than as low-level directives that override safety.

Finally, transparency is important. The robot should provide clear indications of its current state and intentions: where it is about to move, which phase of workflow it is in, how confident it is about the material behavior. Visual signals, sounds, and interface dashboards can all contribute. In effect, the robot becomes a studio colleague whose competence is matched by its predictability and respect for shared space. The skill stack is therefore as much a social technology as a control system, mediating between the raw power of the body stack and the nuanced demands of artistic collaboration.

6. The Art Stack: Style, Aesthetics, and Creative Judgment

If the body and skill stacks let the robot stand in a workshop and competently move tools through stone, the art stack determines *what* it should carve and *when* a given state of the work is good enough, too crude, or overworked. This layer is where historical style, anatomy, composition, and personal taste are encoded in computational form. In humans, these capacities emerge from years of looking, drawing, modeling, and internalizing traditions. For a Michelangelo-class robot, they must instead be constructed from data and models, and then coupled to the physical process through evaluative and generative loops.

The art stack has several interlocking components. It requires a carefully curated corpus of visual and geometric data, enriched with art-historical annotations that capture more than raw shape. It needs a core set of models, here loosely called “Michelangelo-Net,” that can both generate plausible new forms and discriminate between stronger and weaker solutions. It benefits from multi-scale critics that can evaluate work at the level of local tool marks, intermediate surface passages, and overall pose. It must manage the tension between staying

true to a recognizable style and exploring sufficient variation to avoid mere copying, all while respecting the structural demands of marble. Finally, it should be designed for interaction with human artists and historians, who provide guidance, correction, and co-authorship rather than being displaced.

6.1 Data foundations: scans, photographs, and art-historical annotation

Any style-aware system begins with data. For Michelangelo-class sculpting, this data encompasses three main modalities: three-dimensional representations of works, two-dimensional imagery under varied lighting and viewpoints, and textual or symbolic annotations from scholarship and practice.

High-resolution 3D scans of Michelangelo's sculptures and related works form the backbone of geometric understanding. These scans capture not only overall proportions but also small variations in curvature, undercuts, and tool marks. They can be processed into meshes, point clouds, curvature maps, and signed distance fields, each useful for different modeling tasks. The dataset should cover a range of scales, from monumental statues to reliefs and smaller pieces, and also include works by contemporaries and later sculptors in the same tradition to situate Michelangelo's style within a broader manifold.

Photographs provide complementary information. Multi-view image sets under controlled lighting reveal how surfaces respond to light, how tool marks catch shadows, and how polished and rough areas differ visually. Historical photographs and documentary images add context: the sculpture in situ, the interplay with architecture, and the patina acquired over time. By aligning images with scans, the system can learn associations between geometric features and their visual appearance, which is crucial if human feedback will largely be given on images rather than raw geometry.

Art-historical annotations tie these data to meanings and patterns that raw pixels and points do not encode. Annotations might include labels for contrapposto, anatomical features, degrees of non-finito, symbolic motifs, and stylistic phases. They can highlight recurring compositional strategies, such as the twist of torsos, the distribution of weight, or the way drapery frames limbs. They can mark regions where scholars note particular tensions or innovations. These labels transform the dataset from a mere collection of shapes into a structured corpus that reflects a tradition and its internal logic.

6.2 A "Michelangelo-Net" for generative and discriminative style modeling

On top of this data foundation, the core models of the art stack must perform two broad tasks: generation and discrimination. Generation asks, "Given a prompt or constraint, propose a new sculpture in the style of Michelangelo." Discrimination asks, "Given a candidate sculpture or intermediate state, how well does it align with learned stylistic and aesthetic principles?" A

“Michelangelo-Net” is not a single monolithic network but a family of models that collectively address these tasks.

The generative component might be built around a latent space of forms learned from scans and their deformations. Each point in this space corresponds to a plausible configuration of masses, poses, and degrees of finish. Conditioning variables—such as subject, emotional tone, or intended scale—can steer sampling. The model should be able to interpolate between known works while also extrapolating into less explored regions that still feel coherent. Constraints from anatomy, balance, and stone limitations can be incorporated through additional regularizers or auxiliary networks that estimate feasibility.

The discriminative component serves as the robot’s internal critic. It receives as input either full 3D proposals or rendered views of the current state and outputs scores or judgments along relevant dimensions: stylistic consistency, anatomical plausibility, expressive power, and alignment with historical patterns. Different discriminators might be specialized for different tasks: one sensitive to contrapposto and global composition, another focused on anatomical fidelity, another on the distribution and character of tool marks.

Importantly, generation and discrimination should form a closed loop. Candidate designs from the generative model are evaluated by discriminators, which in turn provide gradients or preference signals that refine the generator. This can happen both offline, during pre-training on the static corpus, and online, as the robot works on a particular piece and encounters unforeseen constraints. Over time, Michelangelo-Net evolves a nuanced sense of “what works” within this sculptural language.

6.3 Multi-scale aesthetic critics: from local tool marks to global pose

Aesthetic judgment in sculpture is inherently multi-scale. Viewers move closer and farther, shifting attention from the sweep of a figure’s outline against space to the delicate transition between a tendon and the skin above it. The robot’s evaluative mechanisms must mirror this multi-scale sensitivity if it is to approach human judgments.

At the local scale, critics focus on patches of the surface: small regions of stone where tool marks, micro-curvature, and transitions between adjacent planes define the visual and tactile impression. A local critic might evaluate whether a tooth chisel’s grooves are appropriately spaced and oriented, whether a rasped area has been over-smoothed, or whether a thin ridge is crisp enough to catch light without becoming fragile. Training such a critic involves cropping patches from scans and images, labeling them by quality or style, and learning mappings from local geometry and texture to scores.

At an intermediate scale, critics consider larger passages: an entire muscle group, a fold of drapery, or a section of hair. Here the focus is on how parts relate to each other: the swelling of

a biceps relative to the compression of the forearm, the way a fold of cloth resolves into a sharp edge, or how patterns of rough and smooth areas alternate to create rhythm. These critics must understand anatomy and material behavior more deeply, as they are implicitly evaluating function as well as form.

At the global scale, critics evaluate the overall pose and composition. They assess weight distribution, lines of action, symmetry and asymmetry, and the relationship between the figure and the block. They are sensitive to whether a pose reads as stable, dynamic, or awkward; whether gesture aligns with intended emotion; and whether the silhouette is clear and compelling from key viewpoints. Training them involves exposing the model to many examples of successful and unsuccessful compositions, both within and beyond the Michelangelo corpus.

In practice, these critics operate together. After each significant carving phase, the system can generate a hierarchy of evaluations: local maps highlighting rough patches or inconsistent tool use, intermediate annotations of anatomical or drapery issues, and global judgments about pose and expression. The robot then uses these to prioritize where to work next, much as a human sculptor might scan a piece for "trouble spots" and broader compositional imbalances.

6.4 Balancing style fidelity, diversity, and structural sanity

One of the central challenges for the art stack is to balance three imperatives that can come into tension: staying faithful to Michelangelo's style, generating diverse and original works, and maintaining structural sanity in stone. Each imperative is necessary; none can be pursued without regard for the others.

Style fidelity ensures that outputs are recognizably "akin to Michelangelo." It pushes the generative models toward regions of form space that share characteristic anatomy, poses, surface treatments, and figure–block relationships. Left unchecked, however, this pressure can lead to conservative reproduction: slight variations on a few canonical pieces, or repetitive combinations of familiar motifs.

Diversity counters that tendency by encouraging exploration. It pushes the system to try new subject matter, unusual poses, or novel degrees of non-finito within the learned language. This might involve sampling from less dense regions of the latent space, applying controlled perturbations to conventional compositions, or explicitly optimizing for variety across a series of works. Without diversity, the robot would be a copier; with too much, it might drift into outputs that feel off-style or incoherent.

Structural sanity grounds both in the realities of marble. A composition that is stylistically interesting but structurally impossible is not useful. This constraint must be enforced at multiple stages: the generative model should avoid proposing limbs that are too thin or balanced in precarious ways; the discrimination layer should penalize designs that fail strength or stability

estimates; and the control stack must have the authority to simplify or adjust features that prove unsafe as carving proceeds.

Balancing these imperatives can be framed as a multi-objective optimization problem. During design, the system might search for sculptures that maximize a weighted combination of style fidelity and diversity subject to structural constraints. Weights could be adjusted depending on context: a commission seeking a highly traditional work might emphasize fidelity; an experimental project might prioritize diversity within safety bounds. Over time, feedback from human collaborators and public reception could further calibrate these trade-offs.

6.5 Interfaces for human critique, correction, and co-authorship

Even a highly sophisticated art stack benefits from human participation. Sculptors, art historians, and curators bring perspectives and sensitivities that are difficult to capture fully in data and models. Designing interfaces that allow these experts to critique, correct, and co-create with the robot is therefore essential.

At the simplest level, humans can review renderings or physical maquettes of proposed designs and provide coarse feedback: approving, rejecting, or ranking alternatives. This feedback can be used to fine-tune discriminative models, aligning them more closely with expert taste. More detailed interactions might involve annotating specific regions: circling a shoulder that feels too stiff, drawing an alternative line of action through the body, or marking an area where more roughness should be retained.

The interface can also operate in real time as the robot carves. A sculptor might pause the process, walk around the piece, and point out regions where the robot should focus next. Gesture recognition, laser pointers, or augmented reality overlays can translate these indications into constraints for the planning system. Textual or spoken instructions, such as “soften the transition here” or “leave this section rough,” can be interpreted in light of the art stack’s learned concepts, allowing for high-level control rather than micromanagement.

For deeper co-authorship, the system can support collaborative design sessions. Human participants and the generative model can interactively modify a sculpture in virtual space, with the robot later realizing the agreed design in stone. Alternatively, the robot might propose partial forms or emergent patterns during carving that humans then respond to, creating a dialogue between machine-driven and human-driven decisions. In each case, clear provenance records can document who or what contributed which aspects of the final work, preserving accountability and authorship.

These interfaces serve not only functional but also cultural purposes. They make the robot legible as a tool and partner rather than an opaque autonomous agent. They allow existing communities of practice to adapt their workflows incrementally, integrating the robot as an

advanced instrument rather than as a disruptive replacement. In this sense, the art stack is not just a matter of internal models but of how those models are exposed and shaped through ongoing human–robot collaboration.

7. End-to-End Workflow from Prompt or Maquette to Statue

The three stacks described so far come together in an end-to-end workflow that begins with some form of input and ends with a finished (or deliberately unfinished) statue. This workflow is not a single monolithic algorithm but a layered, cyclical process in which design, evaluation, and execution feed back into each other. It must accommodate very different starting points: a commission to reproduce an existing work, a guided collaboration with a human sculptor, or an open-ended creative brief phrased in natural language.

From the robot's perspective, each project unfolds as a dialogue between internal models, the physical stone, and human interlocutors. An initial design is proposed, refined in virtual space, and then progressively realized in marble through sequences of roughing, blocking, and finishing passes. After each phase, scans and evaluations inform the next decisions, including possible revisions of the design itself. Along the way, choices about non-finito and emergent forms are made explicit parameters of the process rather than accidents or afterthoughts.

This section traces that workflow across five stages. It begins with input modes: copying, guided design, and creative prompts. It then considers how roughing passes are planned with structural safety in mind. Next, it examines iterative scan–evaluate–carve cycles and the criteria by which the system decides that a particular phase or the entire project is complete. It then treats non-finito and controlled emergence as design decisions woven throughout the process. Finally, it sketches hypothetical projects to make these ideas concrete.

7.1 Input modes: copying, guided design, and creative prompts

A Michelangelo-class robot must handle at least three distinct input modes, each with its own implications for the art and skill stacks.

In copying mode, the goal is to reproduce an existing object as faithfully as possible. The input might be a high-resolution 3D scan of a canonical sculpture or a maquette created by a human sculptor. The art stack in this mode does relatively little generative work; its primary responsibility is to interpret the scan into a target form, perhaps smoothing noise, inferring occluded regions, and aligning the model within the chosen block of marble. The robot then treats the target as a fixed reference against which deviations are measured. Even here, however, the style models remain relevant, as they can help the system distinguish between

meaningful features and scanning artifacts, and guide choices about how to realize ambiguous details in ways consistent with the underlying tradition.

In guided design mode, the robot works from partial inputs and ongoing human direction. A sculptor might supply a clay maquette that captures the overall gesture and massing but leaves details unresolved. Alternatively, sketches and annotations can define proportions, key anatomical landmarks, and degrees of finish. The generative side of the art stack uses these as constraints, filling in missing information with proposals that adhere to Michelangelo-esque principles. Human collaborators review and adjust these proposals, perhaps in a virtual environment where they can rotate, scale, and markup the digital model. The outcome is a co-designed target form, neither purely machine-generated nor purely human-drafted.

In creative prompt mode, the input is even more abstract: a text description such as “a figure struggling upward out of the stone, echoing the Prisoners but with a calmer, contemplative face,” optionally accompanied by reference images or mood boards. Here the generative model plays its largest role. It maps the prompt into the latent space of forms, samples candidate designs, and filters them through discriminators sensitive to style, anatomy, and feasibility. The best candidates are rendered for human review, refined through feedback, and eventually selected as targets. In some cases the process may remain open-ended, with the robot beginning to carve a design that is intentionally under-specified, leaving room for emergent decisions later.

Across all three modes, the end of the input stage is the same: a target representation of the figure and its intended relationship to the block, including positional constraints, finish levels, and any explicit non-finito choices. This representation anchors subsequent planning.

7.2 Planning roughing passes with structural safety margins

Once a target form has been defined and positioned inside a chosen block, the next step is to plan roughing passes. Roughing aims to remove large volumes of stone efficiently while preserving enough material around the target surfaces to allow for later corrections and emergent decisions. It is here that structural considerations first enter the workflow in a systematic way.

The planning process begins by computing the volumetric difference between the raw block and the target shape. This difference can be partitioned into layers or shells of material that will be removed in stages. For each region, the system determines an initial safety margin: an offset distance between the roughing surface and the target surface. Thick limbs and torsos may permit smaller margins; thin wrists, ankles, or extended fingers require more conservative buffers. The structural models estimate how much material is needed to support these delicate features during later stages, and the planning system adjusts roughing regions accordingly.

Next, the system considers the order in which material should be removed. Removing mass from one side of a thin column too early may weaken it before adjacent regions are carved; conversely, leaving large overhangs until the end may create unnecessary difficulties. The planner uses structural heuristics or explicit stress simulations to identify sequences that maintain support while still exposing the figure. It may, for example, rough down the backside of a block while preserving a thick “spine” of stone that protects the front-facing figure until blocking is well underway.

Tool selection and strike patterns are then determined. Roughing passes might be executed with heavy point chisels and tooth chisels, working from broad planes inward. The planner must ensure that tool access is available from the chosen directions and that arm configurations can deliver the required forces without compromising stability. It generates tool-path patterns for each region, including stroke spacing and overlap, and associates them with safety parameters: maximum permissible depth per pass, maximum impact energy, and conditions under which the system should pause to reassess.

The result of this stage is a roughing plan: a structured schedule of actions that progressively removes bulk while keeping the stone within safe bounds. This plan is not static; it will be adapted as the robot observes how the actual block responds. But it provides an initial roadmap that reconciles efficiency with caution.

7.3 Iterative scan–evaluate–carve cycles and convergence criteria

Roughing, blocking, and finishing all proceed through iterative cycles in which the robot carves, scans, evaluates, and replans. These cycles are the practical embodiment of the closed-loop architecture connecting the body, skill, and art stacks.

A typical cycle begins with a set of planned strokes for a particular region and phase. The robot executes these strokes, monitoring forces and vibrations to detect anomalies. After a batch of actions, it withdraws tools, clears chips if necessary, and performs a scan. Depending on the phase and required resolution, this might be a fast depth capture of the whole piece, a focused inspection of the worked area, or a sequence of grazing-light photographs.

The scanned data is registered against the existing model of the sculpture, producing an updated representation of its current state. Geometric differencing then reveals how much material has been removed relative to both the previous state and the target form. This information feeds into multiple evaluators. Structural modules check that safety margins are still intact and that no unexpected overcuts or thin spots have appeared. Aesthetic critics assess the local and global effects of the recent changes, updating their estimates of how closely the piece aligns with stylistic and compositional goals.

Based on these evaluations, the planning system updates its actions. If removal is proceeding more quickly than predicted, it may increase safety margins or reduce strike energy. If roughing is too conservative and large material buffers remain that obstruct visibility or later shaping, it may adjust passes to be more aggressive in safe regions. In blocking and finishing phases, the system may also alter local design targets slightly, for example refining the path of a tendon or the angle of a fold, in response to emergent qualities that the critics rate highly.

Convergence criteria determine when a region, phase, or the entire project is complete. For roughing, convergence might mean that all regions are within a predefined distance of the target surfaces and that structural constraints are satisfied. For blocking, criteria might include alignment of key anatomical landmarks within specified tolerances and satisfactory scores from pose and composition critics. For finishing, convergence may be tied to local measures of surface smoothness where desired, appropriate distribution of tool marks, and global aesthetic scores above threshold, combined with a human approval step.

Importantly, convergence is not purely numeric. Human collaborators can override or adjust criteria, deciding that a piece is “done” earlier than the system would conclude, or insisting on further refinement even if internal metrics suggest sufficiency. The workflow must include explicit checkpoints at which such decisions can be made.

7.4 Handling non-finito and controlled emergence as design choices

Non-finito and emergence from stone are among the most characteristic features of Michelangelo’s work, and a credible robot must be able to handle them not as byproducts but as intentional design variables. This requires that degrees of finish be encoded into the target representation and carried through into planning and evaluation.

One way to formalize this is to define a finish field over the surface of the intended sculpture and surrounding block. Each point in this field encodes the desired level of resolution and polish: from raw block, through rough chisel-work, to medium refinement, up to high polish. The field can be designed manually by a sculptor, inferred from references, or generated by the art stack as part of a non-finito style. It may vary smoothly across the figure, with polished faces and hands grading into rough torsos and completely unworked stone.

During planning, this finish field modulates phase transitions and tool usage. In regions marked for low finish, the system may stop after blocking, leaving the figure partially embedded. Safety margins there might remain larger, and roughing marks may be preserved intentionally. In high-finish regions, more aggressive goals are set for finishing phases, with tighter tolerances and more time allocated to subtle surface adjustments.

Emergence becomes a process of revealing the figure along paths suggested by the finish field and by dynamic evaluations. As carving progresses, the art stack may identify emergent forms

or textures in areas originally intended as rough, and propose adjustments to the finish plan: perhaps a particularly evocative set of chisel marks near a shoulder is preserved rather than smoothed, shifting the boundary between finished and unfinished. Human collaborators can also intervene, redrawing finish contours in response to what they see.

The system's critics must be aware of these intended degrees of finish. They should not penalize rough areas for failing to match the precision of polished ones, but instead evaluate them according to appropriate criteria: coherence of tool-mark patterns, harmony with adjacent finished regions, and clarity of the figure–block relationship. Convergence criteria likewise become local and field-dependent. A region flagged as raw block is "done" once structural constraints are satisfied and transitions to neighboring states are aesthetically satisfactory, even if large differences remain between the actual and fully articulated target geometry.

In this way, non-finito is woven into the workflow from the outset, and emergence is treated as a controllable, revisable property rather than a binary of finished versus unfinished.

7.5 Examples of hypothetical projects and process traces

To make the workflow more concrete, consider three hypothetical projects that illustrate different input modes and uses of non-finito.

In the first project, a museum commissions a faithful reproduction of a well-known Renaissance statue for display while the original undergoes conservation. The input is a detailed 3D scan of the original. The robot's art stack cleans and slightly regularizes this scan, ensuring watertight geometry and removing scanning noise, but otherwise preserves its form. The design is positioned within a new marble block, with finish levels set uniformly high except for the plinth.

Roughing passes are planned conservatively, with generous safety margins, and executed around the clock in a secure workspace. Iterative scan–evaluate–carve cycles quickly converge on a blocked-out figure that closely resembles the original. In finishing, local critics pay special attention to surface transitions, while human experts periodically examine the work in person and via high-resolution imagery. Their feedback tunes the robot's assessment of when tool marks are close enough to the original's character. Convergence is declared when geometric deviations fall within the agreed limits and expert viewers deem the copy visually and tactilely acceptable.

In the second project, a contemporary sculptor collaborates with the robot to create a new "Prisoner" figure, inspired by Michelangelo's series but not tied to any specific existing work. The sculptor provides a small clay maquette expressing the desired twist and gesture, along with a sketch indicating where the figure should emerge from and disappear into the block. The robot scans the maquette, fits it into its style manifold, and generates several candidate full-

scale designs that extrapolate the maquette's gesture into more detailed anatomy and block integration, each with its own finish field.

Together, sculptor and robot select and edit one design, adjusting the degree of emergence along different edges of the figure and rebalancing weight distribution to make the pose more daring. The roughing plan explicitly preserves large expanses of block on one side, while aggressively freeing the opposite flank. As carving progresses, the art stack notices particularly striking tool marks on the partially freed arm and proposes leaving them less refined. The sculptor concurs, updating the finish field. The final work exhibits a deliberate tension: a highly articulated head, shoulder, and hand struggling out of rough stone, with the rest of the torso dissolving back into the block. The process trace reveals multiple instances where emergent qualities in the stone influenced mid-course decisions.

In the third project, a cultural institution commissions a more experimental piece: a multi-figure composition loosely inspired by a Pietà, but with abstracted forms and a higher degree of non-finito. The input is mainly textual, supplemented by reference images from different historical periods. The generative model proposes clusters of forms that suggest cradling and support without literal figuration. Art historians and curators review these proposals, selecting one that balances recognizability with abstraction.

Because the composition involves multiple figures and complex interpenetrating forms, structural constraints dominate early planning. The system positions the design in a block with supporting masses that will remain largely uncarved, using them as both literal and symbolic foundations. Roughing removes material only from corridors where the critics agree that negative space is essential, leaving much of the block intact. Blocking focuses on carving strong, readable silhouettes and key gestures rather than detailed anatomy.

Throughout, human collaborators frequently adjust the finish field, deciding which ambiguous forms should remain suggested rather than resolved. The final piece has broad planes of untouched marble intersected by highly charged, semi-abstracted forms that hint at bodies without fully emerging. Viewers can still see traces of the original block, its grain and flaws visible where the robot chose not to intervene. The process trace for this project reads almost like a dialogue: prompt, proposal, carving, surprise, revision, each stage mediated by the art stack and negotiated with human partners.

Across these examples, the same core workflow recurs: inputs of varying specificity lead to target representations with associated finish fields; roughing plans are crafted to respect structural and aesthetic constraints; iterative scan–evaluate–carve cycles refine the work; and decisions about non-finito and emergence are made explicit and revisable. What differs is the balance between reproduction and invention, the tightness of convergence criteria, and the

degree of human involvement. A Michelangelo-class sculpting robot is thus not defined solely by its mechanical capacity, but by its ability to participate flexibly in such workflows, inhabiting a tradition while enabling new expressions within it.

8. Metrics, Benchmarks, and Turing Tests for Sculptural Mastery

Specifying a Michelangelo-class sculpting robot is only meaningful if there are ways to assess whether it has approached that target. Unlike many engineering domains, where performance can be reduced to clear numerical metrics, artistic evaluation involves layered judgments about form, style, expression, and material handling. Nevertheless, it is possible to define structured tests and benchmarks that capture important aspects of sculptural mastery without pretending to compress the entirety of artistic value into a single score.

This section proposes a family of metrics and evaluation protocols. Some focus on perception, asking experts and lay viewers to compare robot-made works to human-made ones. Others analyze geometry, curvature, and surface statistics to quantify how closely a piece follows or extends a stylistic tradition. Structural robustness and long-term stability introduce engineering criteria that any stone sculpture must meet, regardless of authorship. Originality is considered within the context of a learned stylistic manifold, balancing continuity with novelty. Finally, the section sketches how these ideas might be formalized into standardized benchmarks for robotic fine-art practice, analogous to datasets and leaderboards in other areas of machine learning and robotics.

8.1 Perceptual indistinguishability tests with experts and lay viewers

The most intuitively compelling test of sculptural mastery is perceptual: when confronted with a robot-made work alongside human-made ones, how do viewers respond? Perceptual indistinguishability tests borrow the spirit of the classic Turing test, but are adapted to the visual, tactile, and spatial modalities of sculpture.

One protocol invites mixed groups of experts and lay viewers into a gallery where a number of sculptures are displayed. Some are historical originals or high-quality human-made reinterpretations in a Michelangelo-adjacent style; some are created by the robot. Labels are removed or anonymized. Participants are asked to classify each piece along several dimensions: whether they believe it to be old or new, human-made or machine-made, and, in the case of new works, how close they feel it is to the spirit of Michelangelo. They may also be asked to rate each piece for qualities such as anatomical force, expressive power, and cohesion of figure and block.

The results can be aggregated to produce confusion matrices and preference rankings. A robot approaches Michelangelo-class performance when its works are misclassified as human-made at high rates, and when, among new works, they rank comparably to human sculptors who work seriously within that tradition. Importantly, the test should be run under multiple viewing conditions. Photographic comparisons on screens emphasize outline and lighting, while in-person encounters highlight scale, material, and surface. Allowing viewers to walk around and, in some cases, touch the work brings out aspects of spatial composition and tactility that images cannot capture.

Experts and lay viewers play complementary roles. Experts bring trained sensitivity and familiarity with historical works; they are more likely to detect subtle deviations in anatomy, surface handling, or stylistic coherence. Lay viewers contribute a sense of broader communicative impact: whether the work moves them, feels alive, or stands up in the company of canonical pieces. A strong system should fare well with both, but divergences in their responses can be informative, revealing where the robot captures surface traits without deeper structural qualities, or vice versa.

Perceptual tests are necessarily noisy and context-dependent, but they provide a direct measure of how the work is received by its intended audience. They serve as a partial, human-centered Turing test for sculpture.

8.2 Geometric, curvature, and surface-statistics comparisons

Perceptual tests should be complemented by geometric and statistical analyses that can capture similarities and differences more precisely. These analyses treat sculptures as data, extracting measurements from their shapes and surfaces and comparing them across sets of works.

At the most basic level, one can compute distances between corresponding points on robot-made and reference sculptures, assessing how closely a copy reproduces a known work. However, this is not sufficient for creative pieces, and even in copying, simple distance metrics ignore important structural information. More informative are curvature-based descriptors: distributions of principal curvatures across the surface, histograms of mean or Gaussian curvature, and statistics of how curvature varies across anatomically or compositionally meaningful regions.

These descriptors can be compared to reference distributions derived from Michelangelo's corpus and that of related sculptors. A robot whose works have curvature statistics closely aligned with those of the reference set is likely capturing similar balances of planes, bulges, and hollows. Conversely, systematic deviations might indicate tendencies toward over-smoothing, excessive sharpness, or unnatural regularity, all of which are diagnostically useful.

Surface-statistics analyses can go further by examining tool-mark patterns. By detecting and characterizing grooves left by tooth chisels, textures from rasps, and micro-scratches from abrasives, one can build distributions over spacing, orientation, and depth. Comparing these distributions between robot-made and human-made works reveals whether the robot's tool use falls within traditional ranges or exhibits telltale machine regularities. For example, overly uniform spacing or perfectly periodic patterns might betray an algorithmic origin, whereas human-like variation suggests more sophisticated control.

Global measures such as symmetry, balance of mass around a vertical axis, and the distribution of support points can also be quantified. These capture compositional aspects that affect how the sculpture stands and reads in space. While such geometric metrics do not decide aesthetic merit by themselves, they provide objective anchors that complement more subjective judgments and can guide iterative improvements in both design and execution.

8.3 Structural robustness and long-term stability of finished works

No sculpture, however beautiful, can be considered successful if it is structurally unsound. For a robot working in marble at large scales, structural robustness and long-term stability are non-negotiable benchmarks. They ensure that the work can be handled, transported, and displayed without undue risk, and that its features will not quickly degrade or fracture under normal conditions.

Structural benchmarks begin with mechanical analyses. Finite element models can approximate stress distributions within the stone, given its geometry, support conditions, and expected loads. Measures such as maximum stress relative to estimated material strength, safety factors for slender elements, and sensitivity to small perturbations in support can be computed. A Michelangelo-class system should routinely produce designs that meet or exceed agreed safety margins, even for daring poses and thin projections.

Empirical tests provide further validation. For small-scale pieces or test regions, controlled loading can be applied to measure deflection and failure points. Comparing these results to predictions refines the structural models and calibrates safety factors. Over time, a database of designs and their structural performance can inform the art and skill stacks, steering generative models away from risky configurations and helping planning systems avoid undermining critical supports during carving.

Long-term stability involves more than immediate strength. Environmental factors such as temperature fluctuations, humidity, and vibration can contribute to micro-cracking and slow degradation. Surface features like deep undercuts may trap moisture; thin elements may be more susceptible to thermal stress. Conservation science offers criteria and best practices for mitigating these risks, which can be integrated into structural benchmarks. For example, designs

might be evaluated for drainage paths, exposure of vulnerable planes, and suitability for specific installation contexts.

Robot-made works may also be subjected to accelerated aging tests on a small scale: exposing samples to temperature and humidity cycles, salt sprays, or controlled mechanical vibrations, then assessing damage. While entire statues cannot be tested in this way, data from samples carved with similar techniques can inform guidelines and constraints for full-scale projects.

Including structural and conservation-related metrics in the evaluation framework serves two purposes. It aligns the robot's outputs with the practical responsibilities of sculptors and institutions, and it forces the system to integrate aesthetic ambition with physical realism, much as human masters must.

8.4 Assessing originality within a learned stylistic manifold

A central promise of a Michelangelo-class robot is that it can generate not only copies but also original works "akin to Michelangelo." Assessing originality within a specific stylistic manifold requires tools that can distinguish between trivial variation and genuine innovation that still feels anchored in the tradition.

One approach is to represent both historical and robot-made works in a shared latent space learned by the art stack. This space can be constructed from high-level descriptors such as pose parameters, anatomical configurations, figure–block relationships, curvature statistics, and tool-mark patterns. Within this space, clusters and trajectories correspond to particular stylistic tendencies or phases. Michelangelo's own works may occupy a structured region, with neighboring points corresponding to related pieces or variations.

Robot-generated works can then be mapped into the same space. Originality can be conceptualized as distance from known clusters: a piece very close to an existing work may be considered derivative, while one that lies too far away risks leaving the manifold of style entirely. The interesting region is at moderate distances, where the piece is recognizably within the family but introduces new combinations or extensions of existing patterns.

Quantitative measures such as nearest-neighbor distances, cluster membership probabilities, and manifold density estimates can help characterize where a piece lies. However, these should be interpreted in conjunction with human judgments. Expert viewers can be asked to rate how derivative or novel a work feels and whether its novelties are coherent with the style. Over time, these ratings can be used to refine the mapping between latent-space distances and perceived originality.

Another dimension of originality is thematic. Even if formal characteristics remain close to historical models, new subjects or narratives introduced within the same visual language can be

innovative. Detecting and evaluating thematic originality requires annotation and interpretation beyond geometry, involving iconography, symbolism, and context. While more challenging to formalize, it completes the picture: a robot may, for example, use Michelangelo-esque anatomy and non-finito treatment to depict contemporary figures or abstract concepts, stretching the tradition in ways that geometric metrics alone cannot capture.

Ultimately, assessing originality within a learned manifold is about finding a navigable middle ground between slavish imitation and incoherent novelty. The robot's art stack can be trained and tuned to favor that middle ground, and the metrics described here provide ways to monitor and guide its explorations.

8.5 Toward standardized benchmarks for robotic fine-art practice

To move beyond isolated demonstrations toward cumulative progress, the field needs standardized benchmarks for robotic fine-art practice. These benchmarks would play a similar role to datasets and competitions in other areas of robotics and AI, providing shared tasks, evaluation protocols, and reference baselines against which new systems can be compared.

One class of benchmarks could focus on copying tasks. A curated set of reference sculptures, ranging from small fragments to full figures, would be accompanied by high-quality scans and material specifications. Participants would be asked to produce physical replicas in marble or analogous materials under specified conditions. Evaluation would combine geometric and surface-statistics comparisons with blinded perceptual tests, as described earlier. Scores could be reported in terms of accuracy, time, and resource usage.

Another class could target style-consistent invention. Here, the benchmark might define a stylistic domain—such as Michelangelo-esque Renaissance marble figures—and specify constraints such as maximum height, pose categories, or thematic prompts. Systems would generate and carve original pieces within these bounds. Panels of experts and lay viewers would rate the works on style fidelity, originality, expression, and structural soundness. Quantitative metrics from latent-space analyses and surface statistics would supplement these ratings.

Additional benchmarks might focus on subsystems: tool-use benchmarks that assess the quality of chisel marks in controlled test blocks; structural benchmarks that test the ability to produce thin elements within safety margins; or interaction benchmarks that evaluate how effectively a robot responds to human critiques during a timed session.

Standardization would require agreement on materials, tools, environmental conditions, and acceptable levels of human intervention. It would also demand infrastructure: shared scanning facilities, impartial evaluation committees, and secure storage and display of benchmark pieces. Over time, as systems improve, benchmarks can be refreshed with more challenging tasks, such as multi-figure compositions, integration with architectural elements, or cross-style adaptations.

By establishing such benchmarks, the community can situate claims about “Michelangelo-class” robots within a shared evidentiary framework. Engineering advances can be measured not only in terms of abstract capabilities but in the concrete, shared language of sculptural practice. This, in turn, may encourage more meaningful collaborations between roboticists, AI researchers, artists, and institutions, grounded in a common understanding of what counts as success.

9. Broader Implications for Craft, Labor, and Ateliers

A Michelangelo-class sculpting robot is not only a technical artifact. It is also an intervention in living traditions of craft, studio organization, education, conservation, and cultural stewardship. Even if such a system were initially rare and experimental, serious consideration of its implications exposes tensions that reach far beyond marble: the meaning of handwork in an age of embodied automation, the value of apprenticeship when machines can demonstrate techniques endlessly, and the status of artworks when their production is shared between human designers and robotic executants.

This section sketches some of these broader implications. It considers how the roles of human sculptors might shift in studios where robots are present, and how robotic assistants might be used pedagogically rather than as replacements. It explores applications in archiving and conservation, where the ability to reproduce and sympathetically intervene in artworks has particular weight. It then extends the discussion to adjacent crafts such as wood, metal, and architectural ornament, where similar patterns are likely to emerge. Finally, it addresses ethical questions around authorship, value, and cultural heritage that will inevitably arise if machines can convincingly inhabit historic sculptural languages.

9.1 Shifting roles of human sculptors in a robot-populated studio

The introduction of a Michelangelo-class robot into an atelier changes the division of labor. Traditional studios already involve distributed work: masters, assistants, apprentices, and fabricators each contribute to different phases of a piece. Robots initially may be slotted into roles analogous to highly skilled assistants, handling tasks that are physically demanding, repetitive, or logistically challenging.

In such a configuration, human sculptors focus more on conception, early modeling, critical decision points during blocking and finishing, and final adjustments that require subtle, context-dependent judgment. The robot may handle roughing from block to near-figure, perform bulk material removal that would otherwise consume weeks of hammering, and execute certain repetitive surface treatments under human supervision. This shift could extend the working lives of sculptors whose bodies might otherwise be taxed by heavy stonework, enabling them to undertake more ambitious projects with less physical strain.

At the same time, there is a risk that reliance on robots for core manual skills erodes the tacit knowledge that underpins those skills. If fewer apprentices learn to swing a chisel against marble, or to read the stone by feel and sound, the human capacity to sculpt independently may diminish. Studios will need to decide whether to preserve traditional hand-skill training even when machines could do the same tasks more efficiently, treating certain practices as cultural heritage rather than purely as production techniques.

There is also the question of creative authority. A robot capable of style-consistent invention raises the possibility of works proposed primarily by the machine and curated by humans, rather than the other way around. Some artists may embrace this, framing themselves as directors of mixed human–machine ensembles. Others may draw firm boundaries, insisting that robots remain instruments executing human designs. Negotiating these boundaries will shape not only studio workflows but also the identities of sculptors themselves, as they shift from being primarily makers to being hybrid makers–designers–curators of processes.

9.2 New pedagogical uses: robotic assistants as didactic tools

Beyond production, Michelangelo-class robots can serve as powerful didactic tools. Their ability to repeat actions with fine control and to expose internal models makes them ideal for demonstrating techniques and illustrating concepts that are otherwise difficult to convey.

For example, a robot could be programmed to carve a series of small test blocks, each illustrating a specific tool or technique: point chisel roughing at different angles, tooth chisel passes with varied spacing and stroke length, transitions from rough to smooth surfaces, or different treatments of drapery folds. Students could observe these demonstrations from multiple viewpoints, compare results directly, and even experiment with modifying parameters to see the effects. Because the robot can measure forces, vibrations, and material removal with high precision, it can also display quantitative relationships that human practitioners may only intuitively describe.

Robotic assistants could also be used in augmented apprenticeship. A student might rough out a form by hand, then hand the piece to the robot to complete certain phases, or vice versa. The robot's model of the piece and its internal evaluation of surfaces and volumes could be visualized, providing feedback on where the student's work aligns with or diverges from learned stylistic and anatomical norms. This kind of instrumented critique could complement, not replace, human teachers, offering an additional lens through which students can understand their decisions.

In art history and conservation education, robots might recreate historical techniques under controlled conditions: carving replicas of fragments to show how different workshops treated similar motifs, or running comparative experiments on how different chisel geometries affect

tool marks. This can deepen understanding of past practices and support more nuanced attribution and analysis of historical works.

However, pedagogical use must be carefully framed. If students come to rely on robots to correct or complete their work, they may internalize a diminished sense of their own agency. Institutions will need to balance the advantages of high-precision demonstrators with the need to cultivate embodied, personal skill, allowing students to both learn from and, eventually, surpass the machine's capabilities in particular directions.

9.3 Archival and conservation applications of sculpting robots

In archiving and conservation, the ability to carve with high fidelity and stylistic sensitivity has obvious applications. Museums and heritage institutions routinely face the challenge of preserving fragile works, reconstructing damaged pieces, and presenting replicas for display or research. A Michelangelo-class robot can contribute to these tasks in ways that go beyond generic CNC reproduction.

For archival purposes, robots can create physical surrogates of important works that are too fragile to be exhibited or loaned. These surrogates can be used for public display, allowing the original to remain in controlled conditions. Because the robot's art stack understands stylistic nuances, it can produce replicas that capture not just geometry but also the character of tool marks and surface finishes that are crucial to the work's visual impact.

In conservation, robots can assist with interventions that require adding or modifying material. For instance, when replacing a missing section of an architectural frieze or a damaged fragment of a statue, the robot can carve a replacement piece in a matching style, guided by digital reconstructions and the surrounding original material. Its sensitivity to curvature and tool-mark patterns can help ensure that the new section blends harmoniously without falsifying the original beyond accepted thresholds.

Robots can also be used to test conservation treatments on replicas. Before applying a cleaning method, consolidant, or protective coating to an original, conservators might have the robot carve test panels that closely mimic the original's surface, then subject these panels to treatments and accelerated aging. This creates a safer experimental environment and yields more reliable predictions of how interventions will affect the original.

These benefits raise ethical questions. If robots can produce near-indistinguishable replicas, institutions must develop clear policies about labeling, display, and disclosure. Viewers should be able to know when they are seeing an original and when they are seeing a robotic reproduction or a hybrid object, even if their perceptual experience is similar. Otherwise, the line between preservation and substitution may blur in ways that undermine trust.

9.4 Extensions to other crafts: wood, metal, and architectural ornament

While this paper focuses on marble sculpture in a Michelangelo-esque mode, the architecture of a Michelangelo-class robot generalizes to other crafts. Once robots can wield tools with high dexterity and style-aware control, they can be adapted to wood carving, metalworking, and the creation of architectural ornament in a variety of historical idioms.

In wood carving, the challenges shift from brittle fracture to grain-following and the balance between carving and splitting. Style-aware robots could reproduce or extend Baroque or Rococo ornament, folk carving traditions, or contemporary abstract wood sculpture. The sensory and control modules would need to adapt to different force regimes and tool geometries, but the overall structure of body, skill, and art stacks remains relevant.

In metal, especially in repoussé, chasing, or ornamental ironwork, robots can assist with repetitive patterns and precise detailing. Forge work and hot metal introduce additional risks and complexities, but some phases of patterning and finishing are amenable to robotic assistance. Again, style models can guide pattern repetition, variation, and the balance between smooth surfaces and textural accents.

Architectural ornament, from stone tracery to sculpted capitals and friezes, is a particularly promising domain. Here, robots might handle large series of similar elements, introducing controlled variation that keeps ensembles lively while maintaining structural and stylistic coherence. Conservation and restoration of façades and interiors could similarly benefit, with robots carving replacement elements that respect historical patterns while clearly documenting their contemporary origin.

Each extension brings its own cultural concerns. Craftspeople in these domains may worry about displacement or devaluation of handwork, just as sculptors might. Heritage organizations may wrestle with questions about authenticity when new ornament is produced by robots trained on historical styles. The specific dynamics will differ, but the underlying pattern is similar: as robots become capable of inhabiting style-specific craft practices, societies must renegotiate the boundaries between hand, machine, and tradition.

9.5 Ethical questions around authorship, value, and cultural heritage

Finally, the presence of a Michelangelo-class sculpting robot foregrounds ethical questions that cannot be answered by technical specifications alone. Chief among these are issues of authorship, the economic and spiritual value of works produced with robotic assistance, and the stewardship of cultural heritage in an era where stylistic languages can be technically emulated.

Authorship becomes a layered concept. In a simple copying job, authorship of the underlying work remains with the historical sculptor, but the robot and its operators share responsibility for

the reproduction. In co-designed works, where humans and robots exchange proposals and critiques, the question becomes how to apportion creative credit. Some may argue that as long as a human signs off on the final form, authorship remains human; others may see the robot's generative contributions as substantial enough to warrant explicit recognition.

Value is likewise contested. If a robot can produce sculptures that are perceptually and geometrically comparable to those of skilled human sculptors, how should they be priced? Should they command lower prices because they lack the aura of handwork, or should certain robot-made works be valued for their role in advancing the medium? Markets may initially discount robot-assisted pieces, but tastes can evolve, especially if the works achieve genuine expressive or conceptual breakthroughs.

Cultural heritage considerations are even more delicate. Training a robot to sculpt in Michelangelo's style requires digitizing and analyzing his works, which raises questions about rights, permissions, and the responsibilities of institutions that hold those works. Once a style is encoded and can be reproduced indefinitely, there is a risk of trivialization or overuse, flooding markets or public spaces with superficial homages that dull the impact of the originals. Institutions, artists, and policymakers may seek to create norms or regulations that limit certain uses or require contextualization.

There is also the ethical dimension of representation. If robots can convincingly inhabit historical European sculptural languages, similar technologies could be applied to the styles of other cultures and periods. The potential for appropriation, misrepresentation, or commercialization of sacred or culturally specific forms is significant. Engaging with communities whose traditions are being modeled, and respecting their perspectives on what uses are acceptable, will be crucial.

In all of these questions, the presence of a Michelangelo-class robot acts as a catalyst. It does not create entirely new ethical dilemmas so much as intensify existing ones about reproduction, authenticity, and the relationship between technique and meaning. Addressing these dilemmas requires not only engineering caution but also sustained dialogue among artists, scholars, conservators, legal experts, and the public. The technical possibility of sculptural mastery by machines does not determine how societies will choose to integrate that possibility into their cultural and moral landscapes.

10. Conclusion and Future Directions

The thought experiment of a Michelangelo-class humanoid sculptor forces an unusual degree of specificity about what "artistic intelligence" in a physical medium would entail. Instead of stopping at the claim that machines can generate attractive meshes or that industrial arms can

carve stone, this paper has treated sculptural mastery as a layered, embodied capability. It has outlined a body stack of mechanics and sensing tuned to marble, a skill stack that translates high-level intentions into safe, adaptive tool use, and an art stack that internalizes stylistic, anatomical, and compositional knowledge in a way recognizably akin to Michelangelo's practice.

Together, these stacks support an end-to-end workflow from prompt or maquette to statue, traversing roughing, blocking, and finishing phases under continual scan–evaluate–carve cycles. They sit within a broader framework of metrics and benchmarks that seek to evaluate not just accuracy and efficiency but also perceptual reception, stylistic fidelity, originality, and structural integrity. Around this core sit implications for studios, pedagogy, conservation, and cultural heritage, where the presence of such a system would reconfigure existing roles and raise new ethical questions.

The goal has not been to claim that such a robot exists or is just around the corner, but to clarify what building one would demand from robotics, AI, and the humanities. In doing so, several technological gaps, research priorities, and conceptual puzzles come into clearer view.

10.1 Summary of the proposed Michelangelo-class specification

A Michelangelo-class sculpting robot, as specified here, is defined by the convergence of four broad capability domains.

First, it must inhabit the material and tooling ecology of marble sculpture. This means working primarily in marble and related materials, wielding a full repertoire of traditional tools—point, tooth, and flat chisels; rasps and rifflers; drills and abrasives—under impact and sliding regimes comparable to those of human sculptors. The robot's morphology and work envelope must allow it to circle large blocks, reach high and low regions, and lock itself into stable configurations for both heavy roughing and delicate finishing.

Second, it must execute sculptural tasks with high geometric and surface fidelity, as well as stylistically coherent tool-mark signatures. Sub-millimeter accuracy over human-scale figures, controlled distributions of curvature and surface transitions, and human-like variation in tool-mark patterns are all part of this specification. Fidelity is not limited to copying; it includes the ability to inhabit and extend a Michelangelo-adjacent aesthetic vocabulary.

Third, it must be capable of style-consistent invention. Given prompts, maquettes, or guided design constraints, the robot should generate and realize new sculptures that are recognizably in the Michelangelo family without collapsing into direct imitation. This capacity is concentrated in the art stack: a "Michelangelo-Net" that combines generative models of poses, anatomy, and figure–block relationships with discriminative critics that assess style, expression, and structural feasibility.

Fourth, it must operate in closed-loop workflows under explicit evaluation protocols. Iterative scan–evaluate–carve cycles, phase-structured workflows, and multi-scale aesthetic and structural critics ensure that the system is continuously adjusting to the stone and to human feedback. Success is measured not only by geometric closeness to targets but also by how experts and lay viewers perceive the resulting works, how they behave under structural load and environmental exposure, and how original yet coherent they are within a learned stylistic manifold.

Around these capabilities, the specification anticipates integration into real ateliers: human co-authorship, safety and interaction protocols, pedagogical applications, and conservation use cases. The Michelangelo-class robot is thus not a disembodied algorithm but a studio agent embedded in social and material practices.

10.2 Technological gaps and research priorities

Measured against this specification, current technologies fall short in several identifiable ways, which in turn define a research agenda.

On the mechanical and control side, high-force, high-precision humanoid manipulation under repeated impacts remains underdeveloped. Industrial arms and hands typically trade off between stiffness and compliance, and are rarely tuned for hammering tasks where shock absorption, precise timing, and rapid recovery are all essential. Developing series elastic or variable-stiffness actuators, hands with fine-grained force and tactile sensing, and control loops that can handle kilohertz-scale impact dynamics without sacrificing safety is a central priority.

In perception, there is a gap in unified systems that can handle both global 3D reconstruction of large, dusty blocks and close-range inspection of fine surface textures. Existing depth cameras and scanners can capture either large-scale geometry or micro-topology, but integrating them into a robust, dust-tolerant, studio-suitable package with consistent registration over time is still a challenge. Techniques such as grazing-light inspection, photometric stereo, and vibration-based sensing could be combined into multi-modal perception pipelines tailored to stonework.

In modeling and learning, the most significant gaps lie in style-aware generative models for three-dimensional, structurally constrained objects. While two-dimensional style transfer and image generation have advanced rapidly, 3D models that can generate anatomically plausible, physically feasible, art-historically coherent sculptures remain rudimentary. Training such models requires curated datasets of scans, images, and annotations, as well as architectures that treat anatomy, pose, and figure–block relationships as first-class factors rather than incidental outputs.

Multi-scale aesthetic critics likewise pose research challenges. Designing networks that can make meaningful judgments about local tool marks, intermediate surface passages, and global

gestures, and that correlate well with expert human judgments, will require careful dataset construction and ongoing human-in-the-loop calibration. The same is true for structural critics that couple geometry to material behavior in a computationally tractable way.

Finally, evaluation infrastructure is largely absent. Building shared benchmarks, protocols for perceptual tests, standardized structural analyses, and latent-space metrics for originality will demand collaboration between roboticists, computer vision researchers, sculptors, art historians, and conservators. Without such infrastructure, progress will remain anecdotal and difficult to compare across systems.

10.3 From single sculptor-robots to mixed-intelligence ateliers

Even if a single Michelangelo-class robot could be built, a more transformative horizon lies in the development of mixed-intelligence ateliers: studios where multiple robots, humans, and software systems collaborate across projects and media.

In such environments, different robots might specialize in different phases or materials. One system could be optimized for heavy roughing in stone, another for fine finishing and polishing, a third for clay and wax modeling, a fourth for wood or metal. Shared representations—of poses, stylistic manifolds, and figure–block relationships—would allow designs and evaluations to move fluidly between systems, while each robot brings its own physical affordances and limitations to bear.

Humans in these ateliers would inhabit expanded roles. Sculptors might move between hands-on work, virtual sculpting in high-fidelity 3D environments, and curatorial oversight of robot-driven carving. Art historians and conservators might participate directly in design and evaluation processes, using robots to test hypotheses about historical techniques or to explore counterfactual reconstructions. Educators could orchestrate projects where students interact with robots as both tools and interlocutors, learning from their demonstrations and challenging their assumptions.

The art stack, in this context, would evolve into a shared studio intelligence, capable of reasoning across projects and time. It might learn a broader "atelier manifold" encompassing not just Michelangelo's style but the evolving house style of the studio as robots and humans co-create. Over successive projects, this manifold could develop its own internal history: motifs that recur and transform, structural solutions that prove robust, aesthetic experiments that succeed or fail. Robots in the atelier would thus become not just executants but also repositories of institutional memory.

Such mixed-intelligence ateliers raise new organizational and ethical questions. How are decisions made and responsibilities assigned when multiple agents and models contribute? How are shared resources scheduled and prioritized? What forms of documentation and

transparency are needed so that future scholars can understand the roles of various systems in the creation of a given work? Addressing these questions will require as much attention to studio culture and governance as to technical design.

10.4 Open questions at the boundary of engineering and aesthetics

The attempt to specify and eventually build a Michelangelo-class sculpting robot highlights a set of open questions that do not sit neatly within either engineering or aesthetics, but instead span their boundary.

One such question concerns the nature and limits of formalizing style. To what extent can Michelangelo's sculptural language be captured in data and models without flattening it into a set of superficial features? The art stack described here aspires to encode anatomy, pose, figure–block relations, and tool-mark patterns, but there remains a residual sense that something about individual artistic vision resists complete formalization. Whether that "something" is a matter of current technical limitations or an inherent property of artistic practice is an open debate.

Another question is how to understand creativity in mixed human–machine systems. If a robot proposes a novel twist on a canonical pose, and a human sculptor refines and approves it, where does creativity reside? Is it a property of individual agents, of the interaction between them, or of the broader system that includes training data, algorithms, and studio culture? The Michelangelo-class robot is not a solitary genius but a node in a larger network of influences and decisions. Traditional notions of authorship and genius may need to be rethought to accommodate such distributed creativity.

There is also the question of how such technologies might affect our perception of historical works. If robots can approximate Michelangelo's style convincingly, will that deepen appreciation by making his achievements more legible, or will it risk diminishing the aura of the originals by making their effects seem reproducible on demand? Will viewers become more attuned to subtle differences between a master's hand and an encoded style, or will the proliferation of derivatives blur those distinctions?

Ethically, the presence of embodied style-emulating robots amplifies long-standing concerns about reproduction, authenticity, and the commodification of cultural heritage. Decisions about how and when to deploy such systems, in which traditions, and under whose authority, will shape cultural landscapes in ways that cannot be predicted purely from technical capabilities. Engaging affected communities and institutions early and continuously is therefore critical.

Finally, there is the question of what kinds of art we want machines to help bring into the world. Michelangelo's work is only one possible benchmark; many other traditions and aesthetic futures could serve as guides. The specification of a Michelangelo-class robot can be seen as a

test case for how carefully we can think about embedding machines in delicate, meaning-laden practices. Whether we ultimately build such robots or not, the exercise of imagining them with this level of detail can inform broader conversations about embodied AI, craft, and the evolving relationship between technical power and human expression.

In that sense, the Michelangelo-class specification is both an engineering blueprint and a philosophical provocation. It invites us to ask not only “Can we do this?” but also “What would it mean if we did?” and “How should we decide which parts of our artistic inheritance to share with machines?” The answers to those questions will emerge, if at all, not from any one lab or studio, but from a long dialogue among makers, thinkers, and publics who have a stake in what sculptural mastery—human, robotic, or mixed—will look like in the centuries to come.

11. References

Accademia Gallery. (n.d.). *Michelangelo's Prisoners (Slaves).* Galleria dell'Accademia di Firenze. accademia.org+2galleriaaccademiafirenze.it+2

Accademia Florence Tours. (n.d.). *Hall of the Prisoners (Slaves).* Accademia Florence Tours

Bicchi, A. (2000). Hands for dexterous manipulation and robust grasping: A difficult road toward simplicity. *IEEE Transactions on Robotics and Automation, 16*(6), 652–662. centropiaggio.unipi.it+1

Deimel, R., & Brock, O. (2015). A novel type of compliant, underactuated robotic hand for dexterous grasping. In *Robotics: Science and Systems (RSS)*. roboticsproceedings.org

Dudek, G., Owen, G., & Jenkin, M. (1997). Shape representation and recognition from multiscale curvature. *Computer Vision and Image Understanding, 68*(2), 170–189. ScienceDirect

Enayati, N., De Momi, E., & Ferrigno, G. (2020). Variable-impedance and force control for robust compliant interaction. *IFAC-PapersOnLine, 53*(2), 15256–15261. ScienceDirect

Hibbard, H. (1975). *Michelangelo.* New York: Harper & Row. Google Books+1

Hughes, A. (1997). *Michelangelo.* London: Phaidon Press. HathiTrust+2Internet Archive+2

Italian American Herald. (2025, June 13). Michelangelo's "slave" sculptures and the power of the unfinished. italianamericanherald.com

Khatib, O. (1987). A unified approach for motion and force control of robot manipulators: The operational space formulation. *IEEE Journal of Robotics and Automation, 3*(1), 43–53. khatib.stanford.edu+1

Matt Auvinen Studio. (2023, August 30). Reproductions of Michelangelo's "Prisoners" or "Captives." matt auvinen sculpture

Matsuyama, K., Takamatsu, J., & Ikeuchi, K. (2023). Systematic classification of curvature and feature descriptors of 3D shapes. *Entropy, 25*(4), 624. MDPI

Mokhtarian, F., Khalili, N., & Yuen, P. (1999). Multi-scale free-form surface description. Technical report, Centre for Vision, Speech and Signal Processing, University of Surrey. UCL Discovery

Okamura, A. M., Smaby, N., & Cutkosky, M. R. (2000). An overview of dexterous manipulation. In *Proceedings of the IEEE International Conference on Robotics and Automation* (pp. 255–262). BioDex Lab

Park, J. (2006). *Control strategies for robots in contact.* PhD thesis, Stanford University. khatib.stanford.edu

Piazza, C., Grioli, G., Catalano, M. G., & Bicchi, A. (2019). A century of robotic hands. *Annual Review of Control, Robotics, and Autonomous Systems, 2*, 1–32. robots-that-learn.github.io

Robotor. (2022). ROBOTOR: The robot-sculptor that carves Carrara marble. International Federation of Robotics case study. IFR International Federation of Robotics

Robotor. (n.d.). Robot sculpts stone statue identical to Michelangelo's work. robotor.it

Robots Take Over Michelangelo's Workshop. (2022). American Fine Arts Foundry. afafoundry.com

Santello, M., Bianchi, M., Gabiccini, M., et al. (2016). Hand synergies: Integration of robotics and neuroscience for understanding the control of biological and artificial hands. *Physics of Life Reviews, 17*, 1–23. ScienceDirect

Sciolino, E. (2023, December). Can robots replace Michelangelo? *Smithsonian Magazine.* Smithsonian Magazine

Smarthistory. (n.d.). Michelangelo, The Prisoners. Smarthistory

Vasari, G. (1550/1568). *Lives of the most excellent painters, sculptors, and architects.* Various modern editions. Wikipedia+2Project Gutenberg+2

Vasari, G. (1568/2015). *The life of Michelangelo* (Lives of the Artists). Los Angeles: Getty Publications. Getty Museum Store+1

Villani, L., & De Schutter, J. (2016). Force control. In B. Siciliano & O. Khatib (Eds.), *Springer handbook of robotics* (2nd ed., pp. 195–220). Cham: Springer. roboticsproceedings.org

Yang, Y.-L., Lai, Y.-K., Hu, S.-M., & Martin, R. R. (2012). Multi-scale salient features for analyzing 3D shapes. *Journal of Computer Science and Technology, 27*(3), 570–582. Tsinghua University Computer Graphics

Yu, C., Yao, Z., Liu, H., et al. (2022). Dexterous manipulation for multi-fingered robotic hands: A review. *Frontiers in Robotics and AI, 9*, 867182. pmc.ncbi.nlm.nih.gov

The author has published over 4,600 AI-assisted papers at: https://www.researchgate.net/profile/Douglas-Youvan/research . Google Scholar has indexed these preprints, and if you want a particular reference, the AI mode of a Google search (Gemini) has efficiently catalogued these preprints.